# A PHOTO ALBUM OF EARTH: SCHEDULING LANDSAT 7 MISSION DAILY ACTIVITIES

*William Potter and John Gasch*

*National Aeronautics and Space Administration (NASA)*
*Goddard Space Flight Center, Greenbelt, Maryland, USA 20771*
*FAX: 010-301-286-1766, E-mail: William.J.Potter.1@gsfc.nasa.gov*

*Computer Sciences Corporation (CSC)*
*7700 Hubble Drive, Lanham-Seabrook, Maryland, USA 20706*
*FAX: 010-301-794-8260, E-mail: jgasch@csc.com*

**ABSTRACT**

Landsat 7 is a complex, high-capacity remote sensing mission, which has as a goal the allocation of available resources to maximize the coverage and value of the acquired images. This goal necessitates an automated tool for planning and scheduling image acquisitions. This tool must manage the complex scheduling rules, minimize the manual task of planning and scheduling mission activities, and maximize resource usage on board Landsat 7. This paper describes the approach to mission planning and scene acquisition scheduling for the Landsat 7 mission.

**Keywords:** Buffering Pool, Cloud Cover Avoidance, Clustering, Dynamic Priority, Image Scheduler, Landsat 7, MOC, Planning, Scheduling

## 1. INTRODUCTION

Landsat 7 is the newest member of the Landsat family of remote sensing satellites. Landsat 7 will carry on the mission of the aging Landsat 5 spacecraft by acquiring high resolution, multispectral images of the Earth's surface for strategic, environmental, commercial, agricultural, and civil analysis and research. One of the primary mission goals of Landsat 7 is to accumulate and seasonally refresh an archive of global images with full coverage of the Earth's landmasses. This archive will enable further research into seasonal, annual, and long-range trending analysis of such diverse areas as crop yields, deforestation, population growth, and pollution control, to name a few.

The Landsat 7 spacecraft will be operated from the Mission Operations Center (MOC) at NASA's Goddard Space Flight Center in Greenbelt, Maryland. One of the many components of the MOC is the Image Scheduler subsystem. This subsystem is designed to automate mission planning, scheduling, image selection, solid state recorder (SSR) management, ground station scheduling, and command generation. The Image Scheduler automates the complex task of determining the set of scenes to be acquired based on need and/or demand. At the same time, the Image Scheduler considers ground station resource availability, avoidance of foul weather systems, and spacecraft resource conflicts. The Image Scheduler subsystem automatically generates over 95 percent of the spacecraft commands for a typical operational day.

The Landsat 7 spacecraft flies a near-polar Worldwide Reference System (WRS) orbital pattern consisting of 233 orbits every 16 days. Each 98.9-minute orbit is divided into 248 scenes for a total of 57,784 scenes in the global WRS grid. Of these, 13,334 scenes contain enough landmass to be of potential interest to the Landsat 7 mission. Each day the spacecraft will traverse approximately 800 to 850 of these scenes. Duty cycle constraints on the spacecraft instruments limit the total scene acquisition to about 530 scenes per day.

The Image Scheduler acquires images as a result of the following types of requests:

- ***Global archive refresh requests*** for seasonal updates of images worldwide
- ***Special requests*** by United States government agencies for images of high interest
- ***International Ground Station requests*** for the real-time direct downlink of images of their host nations and surrounding regions

The combined scene acquisitions for special requests (SPCs) and global archive refresh requests (GARs) are currently limited to an average of 250 scenes per day. In addition to SPCs and GARs, the mission can acquire 280 more images per day to fulfill international ground station (IGS) scene requests.

The schedulable resources include spacecraft instruments for acquisition and storage of images, and ground stations for downlinking recorded and real-time images. The spacecraft instrumentation includes

- ***Enhanced Thematic Mapper-Plus (ETM+)***. The Image Scheduler must manage the utilization of the ETM+, the scanning mirror imager, within its operating constraints. These constraints include a required warm-up period prior to image acquisition, thermal and power consumption duty cycle budgets, and continuous daytime and nighttime imaging limits.
- ***Solid state recorder (SSR)***. The SSR has a capacity of 46 gigabytes (GB) and can store up to 100 images. The operational capacity is less, typically 90 scenes. This is due to additional data overhead when the stored images are not contiguous scenes.
- Four ***X-band transmitters*** and three gimbaled ***X-band antennae*** for downlinking real-time and stored images.

## 2. THE PLANNING AND SCHEDULING PROBLEM

The Image Scheduler must produce a daily, conflict-free schedule for imaging the “best” 250 scenes for refreshing the global archive. The Image Scheduler equitably distributes the remaining resources to acquire an additional 280 scenes per day to help satisfy requests by international subscribers.

Landsat 7 image scheduling is driven by scene requests. Daily and short-term image requests from science interests and government agencies supplement a database of long-term seasonal imaging requests for populating and refreshing the global archive. Landsat images are maintained at the Earth Resources Observation System (EROS) Data Center, in Sioux Falls, South Dakota. Each image request represents a demand for one scene or contiguous scenes along a single WRS path. The daily requests typically represent demands for images of high interest scenes, i.e., natural disasters.

The Image Scheduler assesses image quality as a tradeoff among mutually conflicting factors:

- Cloud cover avoidance
- Global seasonal coverage
- Scene demand, which includes both latent requests and the priority of the requesting agency
- Image value, based on either scientific significance or potential after-market value

The Image Scheduler must also maximize the utilization of available resources within the bounds of the ETM+ duty cycle constraints.

The planning and scheduling problem domain is co-dependent on the interwoven constraints. Furthermore, it is difficult to predict in advance the precise set of scenes to be acquired due to the dynamics of the cloud cover avoidance rules.

## 3. APPROACH

The scheduling problem for a downlooking, non-gimbaled instrument such as the ETM+, reduced to its simplest abstraction, is a go/no-go decision agent. Of course, there are many instrument and resource constraints that increase the complexity of the problem domain.

### 3.1 ALGORITHM

The Image Scheduler uses a multiple-pass scheduling algorithm. The first pass initially allocates resources for real-time acquisition, recorded image acquisition, and playback of recorded images to a single ground station. Subsequent passes bind SSR data storage and allocate transmitter and antenna resources for simultaneous downlink to multiple ground stations.

The Image Scheduler considers in order each scene to be visited by the spacecraft during the current scheduling period. For each scene, the Image Scheduler determines whether a request for this scene exists. If more than one request for the scene exists, the Image Scheduler chooses the request with the highest dynamic priority. (Dynamic priority is discussed later in this paper). The Image Scheduler then attempts to allocate the ETM+ to acquire the scene. Additionally, if the scene is to be recorded, the Image Scheduler attempts to allocate space on the SSR for this scene. Otherwise, the Image Scheduler allocates a transmitter and antenna for real-time direct downlink of the scene to the ground station requesting the scene.

The scheduling algorithm employs rules based on optimistic resource allocation and look-behind preemption to alter past decisions based on current knowledge. The Image Scheduler optimistically schedules image acquisition as long as required resources are available. Resources are initially allocated to acquire scene images on demand until a resource is exhausted, regardless of the priority of a scene. Once the resource demand exceeds the supply, the scheduler uses a set of rules to implement look-back preemption. The Image Scheduler revisits and possibly reverses earlier decisions in an attempt to reduce demand on the oversubscribed resource. Candidate scenes for this preemption policy are selected from the set of previously scheduled scenes that are holding or have consumed the constrained resource. One or more of these candidate scenes are selected for preemption based on the relative dynamic priorities of the requests for previously scheduled scenes versus the dynamic priority of the current request. For this comparison, the dynamic priority of each previously scheduled scene is additionally adjusted as a function of the relative payback to be achieved as a result of its preemption (discussed further below). For example, if the constraining resource is the ETM+ duty cycle, then the set of potential preemption candidates is limited to all previously scheduled scenes on either end of existing image clusters within the duty cycle window of the ETM+. If the SSR is full, then the set of potential preemption candidates includes only those scenes that were recorded on the SSR and have not yet been played back. One or more previously scheduled scenes may be preempted to free the constraining resource. However, if sufficient quantity of the constraining resource cannot be freed to satisfy the current demand, then any changes made to prior decisions are rolled-back, i.e., the schedule reverts to its original state, and the current scene request is rejected.

Although the algorithm does not look ahead, it is typically used to schedule a time period longer than needed. On Landsat 7, the operational plan is to run the Image Scheduler daily to produce a command load for a 48-hour period. This provides a 24-hour overlap that not only covers contingencies but also enables the Image Scheduler to effectively "look-ahead" by "looking back from the future" as it plans the extra 24 hours.

The algorithm is a linear finite deterministic model. The Image Scheduler is not an optimum scheduler because it does not implement full backtracking to find the most cost effective path to the solution when it must reach back in time to change a prior decision. Preempting a previously scheduled scene is intended to free a specific constraining resource, but releasing that scene may also free other resources making it possible to alter the course of scheduling in the interim. For efficiency the Image Scheduler does not reconsider these images as possibilities. Although it does not necessarily produce the "best" schedule, the Image Scheduler does produce a good, conflict-free schedule of activities within predictable run-time bounds acceptable to mission planning. Typically, the Image Scheduler and the subsequent command load generation will complete within 1 hour. This short run-time enables a quick turn-around to generate a contingency command load in the event a primary load activation is unsuccessful, thereby eliminating the need to build both primary and backup command loads in advance.

## 3.2 CLUSTERING

Images ideally are acquired in clusters of contiguous scenes. Acquiring an image of an isolated scene consumes more resources than acquiring images in a cluster of contiguous scenes because each cluster is enveloped by header and trailer data. The combined resource overhead represented by the header and trailer is equivalent to approximately one extra scene image. Additionally, if the ETM+ is not already active, it must incur a warm-up period, equivalent to three extra scenes, which is deducted from its duty cycle allotment. Therefore, the Image Scheduler grants an extra priority boost to requests for images that are within a cluster of scenes because they use proportionally less resources (SSR space, ETM+ duty cycle, downlink time) than isolated scenes. Elevating the priority of scenes that are within clusters encourages longer clusters, thus increasing the resource efficiency, i.e., image data ÷ (image data + overhead). This cluster priority boost also discourages the breakup of existing clusters when the scheduler later employs preemption rules. Isolated scenes and small clusters are thus more likely candidates for preemption.

The scheduling rules first consider the preemption of isolated scenes and scenes on either end of existing clusters. If these rules fail to free sufficient resources, additional rules permit the splitting of existing long clusters where at least three contiguous scenes are of lower priority than the current scene being considered. This cluster splitting policy addresses a potential limitation of the basic optimistic algorithm. When a cluster is anchored on each end by scenes of high interest, the cluster might contain cloudy images or otherwise low priority scenes because resources were available at the time. In the face of resource limitations, isolated scenes are more likely to be dropped, followed by end scenes of clusters.

### 3.3 DYNAMIC PRIORITY ADJUSTMENT

Landsat 7 image acquisition scheduling is driven by image requests. Each request is assigned a base priority value. When choosing image requests for a specific scene, the Image Scheduler selects the request with the highest priority after adjusting the request priorities up or down by considering the following factors:

- ***Cloud cover***. A request for a scene that is unusually clear is granted a priority boost. Conversely, priority is reduced for a request for a scene with excessive cloud cover. These cloud cover avoidance rules are described in the next section.

- ***Age of request***. A request is granted a priority boost as a function of the number of consecutive missed opportunities. A missed opportunity is either when a scene is not scheduled for acquisition in a prior 16-day cycle, or when an acquired image fails to meet minimum quality standards, as described in the next section. For example, if the last successfully acquired image of a scene was 48 days ago, then a request for this scene is granted a priority boost based on two missed opportunities—from 32 and 16 days ago.

- ***Unfulfilled requests***. Each request has a finite life span, from an effective date to an expiration date. If an image has not been acquired successfully as the life of a request nears its end, its priority is granted an additional boost. This boost is a function of the remaining number of imaging opportunities within the life span of the request, so called the “nearness-to-end” factor. Once an image has been acquired that satisfies a request, then this priority boost factor no longer applies to that request.

- ***Sun angle***. Requests for high latitude scenes are disqualified during the winter season when the Sun angle at 10 a.m. local time exceeds a threshold—typically 85 degrees.

- ***Clustering***. A request’s priority is boosted if the adjacent neighboring scenes are also scheduled for acquisition (as previously described).

The computed dynamic priority is also considered when selecting candidate scene requests for preemption.

### 3.4 CLOUD COVER AVOIDANCE

The mission goal of Landsat 7 is to acquire “substantially cloud-free” images. The Image Scheduler accomplishes this by employing cloud cover prediction data sets from the National Centers for Environmental Prediction (NCEP). These data sets provide global cloud cover predictions for 6, 12, 18, 24, and 36 hours into the future. The predicted cloud cover at a specific scene is determined by grid interpolation between the NCEP data grid and the WRS scene grid. The grid interpolation is followed by a weighted interpolation of the resulting data points in two temporally adjacent prediction data sets surrounding the time when the spacecraft will visit that scene.

The Image Scheduler applies dynamic priority adjustment to a scene request as a linear function of the predicted cloudiness at this scene with respect to its normal seasonal cloud cover. If the cloud cover avoidance policy were based solely on cloud cover, this would have the negative effect of biasing the mission toward arid regions and away from tropical regions of higher interest. Therefore, cloud cover avoidance is based on the relative predicted cloud cover with respect to the seasonal average cloud cover for each scene. Nominal cloud cover is provided in a static lookup table indexed by scene and time of year. If a scene is less cloudy than average, thus there is better visibility, then the request for this scene is awarded a boost in priority. Conversely, more cloud cover than is nominal results in a priority reduction.

The Image Scheduler implements cloud cover avoidance rules as priority adjustment factors rather than absolute filters because the cloud cover data is merely a prediction. This allows other dynamic priority adjustment factors, such as “missed opportunities” and “nearness-to-end”, to overrule the cloud cover prediction if the scene has not been successfully acquired in several prior opportunities. It is better to err on the side of acquiring a predicted cloudy scene than to miss a clear opportunity in the event the cloud cover prediction was wrong.

For scenes acquired and processed by the Landsat 7 Ground Network (LGN) stations, the Image Scheduler receives daily feedback from the Earth Observing System Data and Information System (EOSDIS) Core System (ECS) Distributed Active Archive Center (DAAC) reporting the relative quality of the acquired images. When the actual measured cloud cover in an acquired image is less than the nominal cloud cover for that scene, the Image Scheduler marks that scene as successfully acquired. Otherwise, the scheduler proceeds as though this scene was not acquired. Thus, requests for this scene will be granted the additional missed-opportunity priority boost for the next opportunity in 16 days.

### 3.5 IMAGE BUFFER POOL

The SSR serves a critical role in the scheduling algorithm beyond its practical role as a data buffer for acquired images awaiting downlink to an LGN station. The scheduling algorithm uses the SSR as a scheduling queue, which is a pool of candidate scenes eligible for preemption. If the SSR is full and the current scene needs to be recorded, the Image Scheduler considers all images currently allocated on the recorder as potential candidates for preemption to make room for the current scene. Therefore, it is beneficial to operate the SSR at or near capacity as much as possible. If the SSR is allowed to empty, as a result of excessive downlink bandwidth and/or geographic positioning of LGN stations, then the Image Scheduler would be allowed to acquire cloudy or low-interest images. When the SSR is at or near capacity, the Image Scheduler becomes more selective. An acquisition of a low priority scene is more likely to be preempted by a later more interesting scene before it reaches its turn for playback.

To address this phenomenon, the Image Scheduler implements a “low-water-mark” for the SSR whereby playback to an LGN station is terminated and deferred to a later ground station contact when the SSR space usage falls below this threshold. This policy improves the scheduling algorithm’s decision-making because it allows for a pool of candidate images for preemption. This somewhat decouples the Image Scheduler's behavior from the geographic placement and downlink bandwidth of LGN stations.

Because the Image Scheduler is constantly scheduling and then preempting scenes from the SSR, it defers binding of specific blocks of SSR memory until the Image Scheduler has completed its resource allocation.  A typical scheduling pass for the entire scheduling period is typically 48 hours.

### 3.6  RELATIONAL MODEL

The Image Scheduler is designed using a relational database approach.  The inherent complexity of this scheduling problem is largely embodied in the relational queries. The Image Scheduler models the problem domain in a set of relational database tables. The complexity of the scheduling algorithm is contained within relational queries on these tables.  The following list shows examples of complex rules embodied in database queries:

- Select the highest priority requests for a specific scene
- Select a set of candidate scenes for preemption
- Manage the logical block allocation and address binding on the SSR
- Constrain instrument utilization within the prescribed duty cycle limits
- Allocate transmitter and antenna resources for simultaneous downlink with multiple ground stations

### 3.7  IMPLEMENTATION

The Image Scheduler is implemented in Oracle, a widely available commercial relational database management system. The user interface features of the Image Scheduler are provided by the Mission Operations Planning and Scheduling System (MOPSS), a generic schedule data manager with a convenient graphical timeline.

## 4.  FUTURE PLANS

The Landsat 7 Image Scheduler has already attracted interest from other spacecraft missions. It will be reused for the Earth Orbiter-1 (EO-1) New Millennium Project mission, which has a mission profile nearly identical to that of Landsat 7.  Other missions expressing interest include Astro-E and the Middle Explorer (MIDEX) missions Mapper and Imager.  The Image Scheduler can be adapted to support a variety of spacecraft flying a polar orbit with downlooking data gathering instrumentation.